CMDB 3.0

By Hank Marquis, Director of IT Service Management Consulting, Enterprise Management Associates

ITIL v3 has radically redefined Configuration Management and scrapped the CMDB as we know it. Oh, and this is a really, really good thing!

ITIL v3 is out, and the authors have done the entire industry a very good turn. They have re-defined the term CMDB and dropped Configuration Management as a process.

The CMDB is one of the most desired, misunderstood, over-marketed and most failure prone of all the ITIL. This is partly part because the previous ITIL authors used the term "configuration management" and "database" in its definition.

The confusion from CMDB is directly related to people not understanding configuration management and thinking that CMDB refers to "normal" configuration management, such as managing device configuration files. Then they think the term CMDB actually refers to a single database – "hey, it has database in the acronym!"

This time around, ITIL authors have re-worked the very ideas of asset and Configuration Management, re-defined the concept of the CMDB, demoted Configuration Management as a stand-alone process, and folded the entire collection under the heading of Knowledge Management.

Following, I will explain how these changes could be one of the best things to happen to ITIL and those of us who work with ITIL on a daily basis.

CMDB 2.0

ITIL v2 defines the CMDB as "A database that contains all relevant details of each CI and details of the important relationships between CIs." In my opinion, this single statement has been the cause of most of the grief around CMDB's. Everyone from practitioners to vendors to consultants thought that sentence meant just what it says "a database" – singular. You know, like an Access database or something.

Many an erstwhile ITIL practitioner took the definition to heart and tried to create a monolithic database of information. However, most never actually read the ITIL, so they didn't know that a real CMDB as described by the ITIL (not as defined in the ITIL v2 glossary) is much more of a data federation application than a database.

A CMDB has to federate (connect to and represent) data from other systems – HR systems, Incident and Problem systems, Change Management systems, discovery tools, inventory and audit systems, etc. The actual data stored by the CMDB should be only data not found elsewhere (in definitive and trusted sources) that describes the attributes and relationships of CIs. If a CMDB is a database at all, it is a meta-database – a database that holds data that describes other data. (Whew.)

Aside from not understanding a CMDB, the real issue is that many companies have more than one CMDB – even

though the ITIL v2 made it sound as if there was one and only one CMDB. It is quite common to have a functional CMDB as described by ITIL v2 in many locations within an IT organizations scope of control. Networking usually has one in the form of an NMS. Application development usually has one or two. Often the Service Desk has its own version, and so on.

Many took this reality (myself included) to mean a CMDB then had to be more of a system than a database per se. A "CMDB system" had many databases underneath it. The example I always used (and continue to use) is that of Internet search provider Google.

Google does not replicate all of the websites it indexes – rather it stores bits of information about the websites it discovers and makes this information searchable. My belief, shared with many others, is that the ITIL v2 CMDB is more of an application than a database. And, of course, that is where the controversy began.

CMDB 3.0 TO THE RESCUE!

ITIL v3 defines a CMDB as "A database used to store Configuration Records throughout their Lifecycle. The Configuration Management System maintains one or more CMDBs, and each CMDB stores Attributes of CIs, and Relationships with other CIs."

Notice how the v3 definition allows for multiple CMDBs! Also note how the CMDB is the product of a Configuration Management System (CMS) – not the Configuration Management process. This new definition is dramatic in its implications.

The ITIL v3 describes having multiple CMDBs, which all connect to a single integrated or "core CMDB" (my emphasis.) Not only does this modification of the idea present an easier path to understanding, it sets the stage for sharing information at even higher levels.

ENTER KNOWLEDGE MANAGEMENT

I wrote an article some time ago titled "Knowledge Management – the missing ITIL process." Well, it is no longer missing. The diagram shows how the authors of ITIL v3 have positioned knowledge as the capstone for the once humble CMDB and Configuration Management process. I am going to work from bottom to top to describe this diagram.
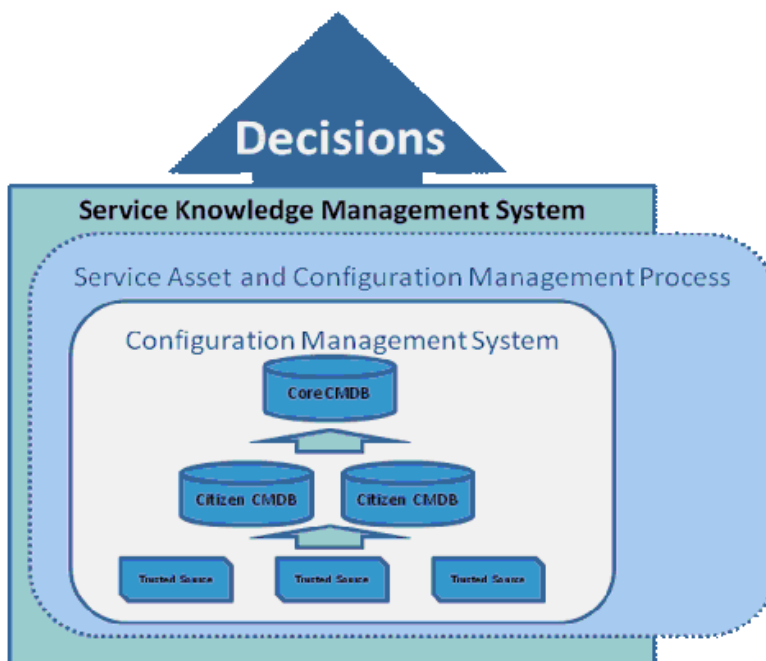

Figure 1. EMA Depiction of Configuration Management Architecture

TRUSTED SOURCES

The various repositories of information that "feed" a CMDB are called data and information sources. EMA refers to these as "trusted sources." A CMDB relies upon its trusted sources.

A CMDB can also be a trusted source for another CMDB. Individual or shared CMDB's, for example from various departments or even across service provider boundaries, are just another source of data and information for a CMDB that may access them.

While the ITIL v3 does not give a name or label to a CMDB acting as a data source to another CMDB, it does apply the label "Integrated CMDB" to the top level CMDB. EMA applies the labels "Citizen CMDB" and "Core CMDB" to unambiguously identify which is which.

CONFIGURATION MANAGEMENT SYSTEM

ITIL v3 introduces the Configuration Management System, or CMS, as the control system or "container" for CI's, management data repositories, CMDB's, and the activities and processes for managing them.

This sounds a lot like the v2 Configuration Management process, and it should. ITIL v3 defines the CMS as "A set of tools and databases that are used to manage an IT Service Provider's Configuration data. The CMS also includes information about Incidents, Problems, Known Errors, Changes and Releases; and may contain data about employees, Suppliers, locations, Business Units, Customers and Users. The CMS includes tools for collecting, storing, managing, updating, and presenting data about all Configuration Items and their Relationships. The CMS is maintained by Configuration Management and is used by all IT Service Management Processes."

Consider the revised ITIL v2 definition (May 1, 2006) for CMDB: "A Database used to manage Configuration Records throughout their Lifecycle. The CMDB records the Attributes of each CI, and Relationships with other CIs. A CMDB may also contain other information linked to CIs, for example Incident, Problem or Change Records. The CMDB is maintained by Configuration Management and is used by all IT Service Management Processes."

To me, it sounds a lot like what a CMDB used to be! Essentially, the old term "CMDB" has been expanded and re-worked to become the CMS. The CMS is "contained" by the new Service Asset and Configuration Management Process.

SERVICE ASSET AND CONFIGURATION MANAGEMENT PROCESS

Under ITIL V2, the financial accounting aspects of managing assets are separate and distinct from managing changes and status of select assets called CI's. ITIL V3 effectively merges all V2 Configuration Management tasks with selected V2 Financial Management tasks to form the new Service Asset and Configuration Management (SACM) Process.

Configuration Management as a stand-alone process is no more. In ITIL v3 Configuration Management is a set of tasks under the SACM process. There is also more clarification about what should be in the CMDB, how CMDBs ought to federate, and so on. As in ITIL v2, a CMDB should reference just those CIs that require Change Management control. Also as in v2, a CMDB should contain the relationships between CIs stored in a CMDB.

SACM has the added responsibility for a broader array of assets, now called "service assets", that also include CIs stored in CMDBs (the reason for the wider box for SACM in the diagram.)

SERVICE KNOWLEDGE MANAGEMENT SYSTEM

Subsuming the CMS (and other informational assets) is a new process called the Service Knowledge Management System, or SKMS.

The SKMS is the set of tools, processes, and databases used to manage knowledge and information. This includes

trusted sources, CMDB's, Configuration Management activities, the CMS, the SACM process and anything else required to turn data into wisdom – the goal of knowledge management and the SKMS.

The SKMS is the capstone on the re-worked configuration management process, and represents the acceptance that knowledge is the fundamental fuel for IT.

SUMMARY

As I have described many times in this column, IT commoditization is unstoppable and it should become the goal of every IT manager and leader. IT commoditization is not a bad thing – just the opposite. IT commoditization means IT has "grown up" and is now so embedded into the lives of people and businesses of all types that it can no longer operate in a haphazard fashion.

All versions of the ITIL have addressed commoditization – ITIL exists because of IT commoditization. Thus, the organization – process, workflow, teamwork, leadership – is more important than any individual piece of technology. How people work requires knowledge, and knowledge management is the fundamental glue that binds an IT organization and transitions it from good to great.

The clarification ITIL V3 brings to these issues is tacit acceptance of these facts, and bodes well for its future and the future of those who choose to focus on it.